

Alternate framework for distributed computing tames Big Data's ever growing costs

February 23 2023

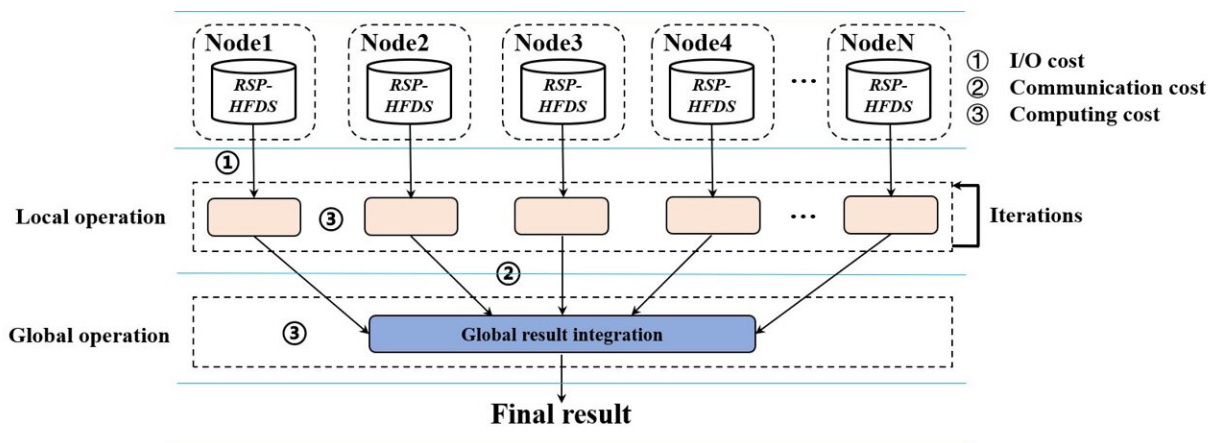


Diagram showing locations of various computational and other costs incurred in the proposed Non-MapReduce framework. Credit: Big Data Mining and Analytics, Tsinghua University Press

The sheer volume of 'Big Data' produced today by various sectors is beginning to overwhelm even the extremely efficient computational techniques developed to sift through all that information. But a new computational framework based on random sampling looks set to finally tame Big Data's ever-growing communication, memory and energy costs into something more manageable.

A paper describing the framework was published in the journal *Big Data Mining and Analytics*.

The amount of data being produced from social networks, business transactions, the 'Internet of Things,' finance, healthcare and beyond has exploded in recent years. This era of so-called Big Data has offered incredible statistical power to discover patterns and deliver insights previously unimaginable. But the volume of Big Data being produced is beginning to hit computational limits.

The scalability of complex algorithms starts to flounder at around a terabyte of data—or one trillion bytes—in a computer cluster or in [cloud computing](#). The New York Stock Exchange for example produces about a terabyte's worth of trade data every day, while Facebook users generate 500 terabytes over that same time.

Distributed computing plays a vital role in the storing, processing and analysis of such big data. This framework deploys a 'divide and conquer' strategy to efficiently and speedily sort through it. This involves the partitioning of a big data file into a number of smaller files called 'data block files.'

These data blocks are stored in a distributed fashion across the many nodes of a cluster of computers. Each of these blocks are then processed in parallel instead of sequentially, radically speeding up processing time. The results from these local nodes are then fed back to a central location and reintegrated, producing a global result.

This divide-and-conquer operation is in turn managed by a distributed file system, which in turn is governed by a programming model. The file system is what divides the big data files, and the programming model divides an algorithm into pieces, which can then run on the data blocks in a distributed fashion.

MapReduce, developed by Google, is the most widely used programming model for distributed computing that runs on clusters and across the cloud. The name comes from its two basic operations. The Map operation is performed on the data block in a node to generate a local result. This is executed on multiple nodes in parallel to achieve the huge speed-up in processing time. The Reduce operation then collates all these local results into a global result.

This latter stage involves a transfer of local results to other master nodes or central node that perform the Reduce operation, and all this data shuffling is extremely costly in terms of communication traffic and memory.

"This enormous communication cost is manageable up to a point," said Xudong Sun, the lead author of the paper and a computer scientist with the College of Computer Science and Software Engineering at Shenzhen University. "If the desired task involves only a single pair of Map and Reduce operations, such as counting the frequency of a word across a large number of web pages, then MapReduce can run extremely efficiently across thousands of nodes over even a gargantuan big-data file."

"But if the desired task involves a series of iterations of the Map and Reduce pairs, then MapReduce becomes very sluggish due to the large communication costs and consequent memory and computing costs," he added.

So the researchers developed a new distributed computing framework they call Non-MapReduce to improve the scalability of cluster computing on big data by reducing these communication and memory costs.

To do so, they depend upon a novel data representation model called

random sample partition, or RSP. This involves a random sampling of a big data file's distributed data blocks instead of a processing of all the distributed data blocks. When a big data file is analyzed, a set of RSP data blocks are randomly selected to be processed and then subsequently integrated at the global level to produce an approximation of what the result would have been had the entire data file been processed.

In this way, the technique works in much the same way as in statistical analysis, [random sampling](#) is used to describe the attributes of a population. The Non-MapReduce's RSP approach is thus a species of what is called 'approximate computing,' an emerging paradigm in computing to achieve greater energy efficiency that delivers only an approximate rather than exact result.

Approximate computing is useful in those situations where a roughly accurate result that is computationally cheaply achieved is sufficient for the task at hand, and superior to a computationally costly effort at trying to deliver a perfectly accurate result.

The Non-MapReduce computing framework will be of considerable benefit for a range of tasks, such as quickly sampling multiple random samples for ensemble machine learning; to directly execute a sequence of algorithms on local random samples without requiring data communication amongst the nodes; and easing the exploration and cleaning up of big data. In addition, the framework saves a significant amount of energy in cloud computing.

The team now hope to apply their Non-MapReduce framework to some major big data platforms and use it for real-world applications. Ultimately, they hope to use it to tackle application problems of analyzing extremely [big data](#) distributed across several data centers.

More information: Xudong Sun et al, Survey of Distributed

Computing Frameworks for Supporting Big Data Analysis, *Big Data Mining and Analytics* (2023). [DOI: 10.26599/BDMA.2022.9020014](https://doi.org/10.26599/BDMA.2022.9020014)

Provided by Tsinghua University Press

Citation: Alternate framework for distributed computing tames Big Data's ever growing costs (2023, February 23) retrieved 27 April 2024 from <https://techxplore.com/news/2023-02-alternate-framework-big.html>

This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.